

WHAT IS CLAIMED IS:

- 1 1. A method, comprising:
 - 2 partitioning a plurality of processing nodes in a storage system into a plurality of logical processing units, wherein the plurality of logical processing units can respond to I/O requests from a host coupled to the storage system;
 - 5 grouping at least two logical processing units, wherein data in a first storage coupled to a first logical processing unit of the least two logical processing units is mirrored by data in a second storage coupled to the second logical processing unit of the at least two logical processing units; and
 - 9 in response to a failure of the first logical processing unit, responding to an I/O request from the host via the second logical processing unit.
- 1 2. The method of claim 1, wherein the storage system has at least two processing nodes, wherein the plurality of logical processing units are distributed across the at least two processing nodes, wherein one processing node includes a plurality of central processing units, and wherein in the event of the failure of the first logical processing unit, the plurality of processing nodes stay operational.
- 1 3. The method of claim 1, wherein an administrative console is coupled to the plurality of processing nodes of the storage system, further comprising:
 - 3 prior to partitioning, processing, at the administrative console, information on processing requirements, memory requirements and host bus adapter requirements for the plurality of logical processing units.
- 1 4. The method of claim 1, wherein one or more partitioning applications are coupled to the plurality of logical processing units, further comprising:
 - 3 in response to grouping the at least two logical processing units, starting initial program load of the first logical processing unit;

5 determining via the one or more partitioning applications an identification of the
6 second logical processing unit grouped with the first logical processed unit; and
7 presenting, by the one or more partitioning applications, common resources to the
8 first and second logical processing units.

1 5. The method of claim 1, further comprising:
2 receiving from the first logical processing unit, a request for memory access of a
3 logical processing unit;
4 determining, by one or more partitioning applications coupled to the plurality of
5 logical processing units, whether the logical processing unit is grouped with the first
6 logical processing unit;
7 if the logical processing unit is grouped with the first logical processing unit, then
8 allowing the memory access of the logical processing unit to the first logical processing
9 unit; and
10 if the logical processing unit is not grouped with the first logical processing unit,
11 then preventing the memory access of the logical processing unit to the first logical
12 processing unit.

1 6. The method of claim 1, further comprising:
2 receiving from the first logical processing unit, a request for a reinitialized
3 program load of a logical processing unit;
4 determining, by one or more partitioning applications coupled to the plurality of
5 logical processing units, whether the logical processing unit is grouped with the first
6 logical processing unit;
7 if the logical processing unit is grouped with the first logical processing unit, then
8 allowing the reinitialized program load of the logical processing unit; and
9 if the logical processing unit is not grouped with the first logical processing unit,
10 then preventing the reinitialized program load of the logical processing unit.

1 7. The method of claim 1, further comprising:

2 receiving a write request from the host to the plurality of processing nodes in the
3 storage system; and

4 writing, by one or more partitioning applications, data corresponding to the write
5 request to the first storage coupled to the first logical processing unit and the second
6 storage coupled to the second logical processing unit.

1 8. The method of claim 1, further comprising:
2 receiving a read request from the host to the plurality of processing nodes in the
3 storage system; and

4 reading, by one or more partitioning applications, data corresponding to the read
5 request from the first storage coupled to the first logical processing unit.

1 9. The method of claim 1, wherein the partitioning and grouping are
2 performed by one or more partitioning applications coupled to the plurality of processing
3 nodes, wherein the one or more partitioning applications comprise a hypervisor
4 application of a redundant system.

1 10. The method of claim 1, wherein the partitioning and grouping further
2 comprises:

3 associating first pool numbers sequentially to a first set of logical processing units
4 included in a first processing node;

5 associating second pool numbers sequentially to a second set of logical processing
6 units included in a second processing node; and

7 assigning pairs of logical processing units with the same associated pool numbers
8 to be partner virtual machines, wherein the partner virtual machines mirror each other and
9 provide redundancy.

1 11. A system, comprising:
2 a storage system;
3 a plurality of processing nodes in the storage system;

4 means for partitioning the plurality of processing nodes in the storage system into
5 a plurality of logical processing units, wherein the plurality of logical processing units can
6 respond to I/O requests from a host coupled to the storage system;

7 means for grouping at least two logical processing units, wherein data in a first
8 storage coupled to a first logical processing unit of the least two logical processing units
9 is mirrored by data in a second storage coupled to the second logical processing unit of
10 the at least two logical processing units; and

11 means for responding to an I/O request from the host via the second logical
12 processing unit in response to a failure of the first logical processing unit.

1 12. The system of claim 11, wherein the storage system has at least two
2 processing nodes, wherein the plurality of logical processing units are distributed across
3 the at least two processing nodes, wherein one processing node includes a plurality of
4 central processing units, and wherein in the event of the failure of the first logical
5 processing unit, the plurality of processing nodes stay operational.

1 13. The system of claim 11, further comprising:
2 an administrative console coupled to the plurality of processing nodes of the
3 storage system; and
4 means for processing, at the administrative console, information on processing
5 requirements, memory requirements and host bus adapter requirements for the plurality of
6 logical processing units.

1 14. The system of claim 11, further comprising:
2 one or more partitioning applications coupled to the plurality of logical
3 processing units;
4 means for starting initial program load of the first logical processing unit, in
5 response to grouping the at least two logical processing units;

6 means for determining via the one or more partitioning applications an
7 identification of the second logical processing unit grouped with the first logical
8 processed unit; and

9 means for presenting, by the one or more partitioning applications, common
10 resources to the first and second logical processing units.

1 15. The system of claim 11, further comprising:

2 means for receiving from the first logical processing unit, a request for memory
3 access of a logical processing unit;

4 means for determining, by one or more partitioning applications coupled to the
5 plurality of logical processing units, whether the logical processing unit is grouped with
6 the first logical processing unit;

7 means for allowing the memory access of the logical processing unit to the first
8 logical processing unit if the logical processing unit is grouped with the first logical
9 processing unit; and

10 means for preventing the memory access of the logical processing unit to the first
11 logical processing unit if the logical processing unit is not grouped with the first logical
12 processing unit.

1 16. The system of claim 11, further comprising:

2 means for receiving from the first logical processing unit, a request for a
3 reinitialized program load of a logical processing unit;

4 means for determining, by one or more partitioning applications coupled to the
5 plurality of logical processing units, whether the logical processing unit is grouped with
6 the first logical processing unit;

7 means for allowing the reinitialized program load of the logical processing unit; if
8 the logical processing unit is grouped with the first logical processing unit; and

9 means for preventing the reinitialized program load of the logical processing unit
10 if the logical processing unit is not grouped with the first logical processing unit.

1 17. The system of claim 11, further comprising:
2 means for receiving a write request from the host to the plurality of processing
3 nodes in the storage system; and
4 means for writing, by one or more partitioning applications, data corresponding to
5 the write request to the first storage coupled to the first logical processing unit and the
6 second storage coupled to the second logical processing unit.

1 18. The system of claim 11, further comprising:
2 means for receiving a read request from the host to the plurality of processing
3 nodes in the storage system; and
4 means for reading, by one or more partitioning applications, data corresponding
5 to the read request from the first storage coupled to the first logical processing unit.

1 19. The system of claim 11, wherein the partitioning and grouping are
2 performed by one or more partitioning applications coupled to the plurality of processing
3 nodes, wherein the one or more partitioning applications comprise a hypervisor
4 application of a redundant system.

1 20. The system of claim 11, wherein the partitioning and grouping further
2 comprises:
3 means for associating first pool numbers sequentially to a first set of logical
4 processing units included in a first processing node;
5 means for associating second pool numbers sequentially to a second set of logical
6 processing units included in a second processing node; and
7 means for assigning pairs of logical processing units with the same associated
8 pool numbers to be partner virtual machines, wherein the partner virtual machines mirror
9 each other and provide redundancy.

1 21. An article of manufacture, wherein the article of manufacture is capable of
2 causing operations, the operations comprising:

3 partitioning a plurality of processing nodes in a storage system into a plurality of
4 logical processing units, wherein the plurality of logical processing units can respond to
5 I/O requests from a host coupled to the storage system;

6 grouping at least two logical processing units, wherein data in a first storage
7 coupled to a first logical processing unit of the least two logical processing units is
8 mirrored by data in a second storage coupled to the second logical processing unit of the
9 at least two logical processing units; and

10 in response to a failure of the first logical processing unit, responding to an I/O
11 request from the host via the second logical processing unit.

1 22. The article of manufacture of claim 21, wherein the storage system has at
2 least two processing nodes, wherein the plurality of logical processing units are
3 distributed across the at least two processing nodes, wherein one processing node
4 includes a plurality of central processing units, and wherein in the event of the failure of
5 the first logical processing unit, the plurality of processing nodes stay operational.

1 23. The article of manufacture of claim 21, wherein an administrative console
2 is coupled to the plurality of processing nodes of the storage system, the operations
3 further comprising:

4 prior to partitioning, processing, at the administrative console, information on
5 processing requirements, memory requirements and host bus adapter requirements for the
6 plurality of logical processing units.

1 24. The article of manufacture of claim 21, wherein one or more partitioning
2 applications are coupled to the plurality of logical processing units, the operations further
3 comprising:

4 in response to grouping the at least two logical processing units, starting initial
5 program load of the first logical processing unit;

6 determining via the one or more partitioning applications an identification of the
7 second logical processing unit grouped with the first logical processed unit; and

8 presenting, by the one or more partitioning applications, common resources to the
9 first and second logical processing units.

1 25. The article of manufacture of claim 21, the operations further comprising:
2 receiving from the first logical processing unit, a request for memory access of a
3 logical processing unit;

4 determining, by one or more partitioning applications coupled to the plurality of
5 logical processing units, whether the logical processing unit is grouped with the first
6 logical processing unit;

7 if the logical processing unit is grouped with the first logical processing unit, then
8 allowing the memory access of the logical processing unit to the first logical processing
9 unit; and

10 if the logical processing unit is not grouped with the first logical processing unit,
11 then preventing the memory access of the logical processing unit to the first logical
12 processing unit.

1 26. The article of manufacture of claim 21, the operations further comprising:
2 receiving from the first logical processing unit, a request for a reinitialized
3 program load of a logical processing unit;

4 determining, by one or more partitioning applications coupled to the plurality of
5 logical processing units, whether the logical processing unit is grouped with the first
6 logical processing unit;

7 if the logical processing unit is grouped with the first logical processing unit, then
8 allowing the reinitialized program load of the logical processing unit; and

9 if the logical processing unit is not grouped with the first logical processing unit,
10 then preventing the reinitialized program load of the logical processing unit.

1 27. The article of manufacture of claim 21, the operations further comprising:
2 receiving a write request from the host to the plurality of processing nodes in the
3 storage system; and

4 writing, by the one or more partitioning applications, data corresponding to the
5 write request to the first storage coupled to the first logical processing unit and the second
6 storage coupled to the second logical processing unit.

1 28. The article of manufacture of claim 21, the operations further comprising:
2 receiving a read request from the host to the plurality of processing nodes in the
3 storage system; and
4 reading, by one or more partitioning applications, data corresponding to the read
5 request from the first storage coupled to the first logical processing unit.

1 29. The article of manufacture of claim 21, wherein the partitioning and
2 grouping are performed by one or more partitioning applications coupled to the plurality
3 of processing nodes, wherein the one or more partitioning applications comprise a
4 hypervisor application of a redundant system.

1 30. The article of manufacture of claim 21, wherein the partitioning and
2 grouping further comprises:
3 associating first pool numbers sequentially to a first set of logical processing units
4 included in a first processing node;
5 associating second pool numbers sequentially to a second set of logical processing
6 units included in a second processing node; and
7 assigning pairs of logical processing units with the same associated pool numbers
8 to be partner virtual machines, wherein the partner virtual machines mirror each other and
9 provide redundancy.